

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.



(19) **EUROPEAN PATENT APPLICATION**
(11) **EP 1 079 581 A2**

(43) Date of publication:
28.02.2001 Bulletin 2001/09
(21) Application number: 00306813.7
(22) Date of filing: 09.08.2000

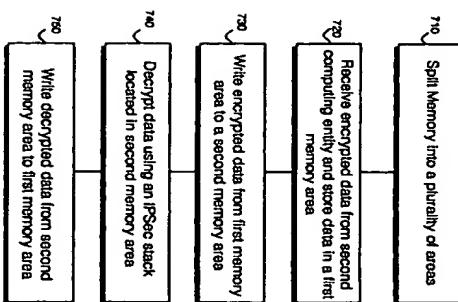
(45) Int. Cl.7: **H04L 29/06, H04L 9/00**

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LU NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 17.08.1999 GB 9919444
(71) Applicant:
Hewlett-Packard Company
Palo Alto, California 94304-1112 (US)

(54) **Robust encryption and decryption of packetized data transferred across communications networks**

(57) A method for improving fault tolerance and data throughput rates in the transfer of data between communicating entities across a virtual private network comprises the steps of: logically dividing a memory into a plurality of areas; receiving encrypted data from said second communicating entity; storing said encrypted data in the first memory area associated with said first communicating entity; writing said encrypted data stored in said first memory area into a second memory area associated with said first communicating entity; decrypting said encrypted data stored in said second memory area; and writing said decrypted data from said second memory area to said first memory area.



1 Description
2 Field of the Invention

1 **EP 1 079 581 A2**
2

[0001] The present invention relates to a method of processing encrypted data packets transmitted between communicating entities across a communications network.

Background to the Invention

[0002] It is known to transfer digital data between communicating entities such as digital computers and the like across computer networks such as wide area networks (WAN), local area networks (LAN) and the Internet. This digital data may represent digitized audio or video signals or increasingly sensitive financial information such as credit card details, banking information and the like.

[0003] Referring to Fig. 1 herein, there is illustrated schematically two such communicating entities 100 and 101 configurable to exchange data across a communications network 102.

[0004] With the development of technologies such as the world wide web (WWW) there has been a rapid increase in the use of networks of computing entities to exchange sensitive information carried as data across the networks. As the amount of sensitive data across transferred across networks has increased there has been increasing pressure to develop and improve the security of transferred data. It is known to improve the security of data transfer between digital computers by encrypting the data using one or a plurality of standard encryption algorithms and transmitting the encrypted data across a network where it may be decrypted and subsequently processed only by those recipients for whom the encrypted data is intended.

[0005] Referring to Fig. 2 herein, there is illustrated schematically a conventional system for encrypting user data and transferring the encrypted data across a network for subsequent decryption. A first digital computer 200 is represented schematically to include the following sub-components:

- A kernel/operating system (OS) 201 which occurs within a logical location or plurality of logical locations within an electronic memory of the first digital computer 200;
- A plurality of user applications 202 — 204 which are stored within a second memory area which is logically distinct from the memory locations storing the operating system, the second memory area also known herein as 'user memory'; and
- A network interface card 205 which is configurable to convert digital data into a form suitable for physical transmission across a network 250.

[0006] The kernel/operating system 201 which occurs within the logical location or plurality of logical locations herein are also known as 'Bump In the Stack (BITS) systems. The Internet protocol security layer 208 accesses a first data base of private key information 209 to encrypt or decrypt data. The private key database 209 contains a name private key data as second private key data base 210 which is associated with a remote second computer 210 which is itself configured to exchange encrypted data with first computer 200 and has a corresponding protocol stack to the protocol stack of the first computer 200. The second computer 210 comprises a kernel/operating system 211, a plurality of applications 212, 213 and a network interface card 215 and has an internet protocol security layer 218. Typi-

		5	EP 1 079 581 A2	4
		6		
	cell, encryption and decryption of user data may be carried out using a standard cipher such as the prior art Data Encryption Standard (DES);			
[0009]	Referring to FIG. 3 herein there is illustrated schematically a conventional unencrypted data packet 300 with an IP header 301 as transmitted across a conventional network. An unencrypted data packet is also illustrated schematically an encrypted data packet 302 having both an IP header 303 and an internet protocol security header 304. Data encryption is carried out by internet protocol security layer 208 which also appends the internal protocol security header data 304 to the encrypted data 305.			
[0010]	Referring to FIG. 4 herein, there is illustrated schematically a conventional encryption system incorporating an internet protocol security layer 400 and associated keying data stored in a key database 401. Inserted into a network protocol stack resident within a kernel/operating system, is an Internet Key Exchange protocol 402 which is designed for negotiating the exchange of keying material, such as contained within key data base 401, between computer entities which wish to exchange data encrypted using the internet protocol security stack 400. In a conventional virtual private network, the internet key exchange protocol resides within a region of user memory, outside of the kernel/operating system.			
[0011]	However, there are significant problems introduced in the prior art when the internet protocol security stack is located within the kernel/operating system. Locating the internet protocol security stack within the network protocol stack necessitates that, in a conventional operating system, the internet protocol security stack must operate in supervisory mode. In the event of a fault occurring during an encryption or decryption operation then the entire operating system may be affected. Malfunctioning of the internet protocol security stack may arise from the internal protocol security stack receiving incorrectly formatted data packets, for example as a result of a deliberate attack by an individual wishing to cause damage to the system. The consequences of receiving an incorrectly formatted data packet are that it may result in a buffer memory overflow or in the overwriting of memory locations reserved for the operating system. This can result in causing the kernel/operating system to crash which consequently causes remaining systems, which depend upon the kernel/operating system, to crash.			
[0012]	Additionally, conventional operating systems are single threaded and non-reentrant and hence are not designed to run on multiple processors. The encryption or decryption of a plurality of data packets can occupy a substantial fraction of total processor resources available to an operating system with the result of a reduction in the fraction of processing resources available to user processes, which thereby affects overall system performance.			
	[0013] It is important to develop more fault tolerant encryption and decryption of user data for transfer across virtual private networks in order to improve the resistance to malicious attack and operating efficiency of networked computers.			
5	Summary of the invention			
	[0014] Specific embodiments and methods according to the present invention aim to improve the tolerance to faults in computer systems configurable to exchange encrypted digital data across virtual private networks by isolating error prone encryption/decryption means from a computer operating system and thereby improve the reliability and security of computer systems.			
10	[0015] According to a first aspect of the present invention there is provided a method of processing encryption data in a computing entity said method comprising the steps of:			
15	assigning a memory means of said computing entity into a plurality of memory areas;			
20	writing said stored data into a second memory area of said plurality of memory areas associated with said first communicating entity;			
25	encrypting said data stored in said second memory area to said first memory area;			
30	writing said encrypted data from said second memory area assigned for use by a kernel of an operating system of said computing entity;			
35	decryption said encrypted data in a first said memory area of said computing entity, said first memory area assigned for use by a kernel of an operating system of said computing entity;			
40	writing said decrypted data stored in said first memory area into a second memory area associated with said computing entity;			
45	decryption said encrypted data stored in said second memory area, and			
50	writing said decrypted data from said second memory area to said first memory area.			
	[0016] Preferably said first memory area is configured to contain code of an operating system program for controlling said computing entity.			
	[0021] Preferably said first memory area is logically distinct from said second memory area, writing said encrypted data from said first memory area to said first memory area.			
	[0022] Preferably said second memory area is partitioned off from said first memory area such that said second memory area is not used for operation of said operating system.			
	[0023] Subsidiarily said redirection means is configured to receive decrypted data from said second memory area.			
	[0024] According to a third aspect of the present invention there is provided a method of processing encrypted data in a computing entity, said computing entity comprising:			
	8 processor; and			
	8 memory means,			
	wherein said memory means is divided into first and second memory areas, wherein said first memory area contains code of an operating system of said computer entity, said method comprising the steps of:			
	receiving an encrypted data packet;			
	processing said data packet according to at least one packetization protocol of said operating system in said first memory area;			
	outputting said data packet to said second memory area;			
	processing said data packet according to a decryption algorithm in said second memory area, and			
	returning said processed data packet to said operating system in said first memory area.			
	[0017] Preferably said first memory area is partitioned from said second memory area such that said second memory area is not used for operation of said computing entity.			
	[0018] Preferably said second memory area is partitioned off from said first memory area such that said second memory area is not used for operation of said computing entity.			
	[0019] According to a second aspect of the present invention there is provided a method of processing encryption data in a computing entity said method comprising the steps of:			
	computing entity into a plurality of memory areas;			
	storing decrypted data in a first said memory area			
	[0025] According to a fourth aspect of the present invention there is provided a digital computer configura-			

[0030] Preferably, said encryption means logically located within said second memory area comprises:

a plurality of Internet protocol security stacks; and

a plurality of data bases configurable to contain key data for encrypting said data, wherein each data base of said plurality of databases contains a same key data.

[0031] Said plurality of Internet protocol security stacks may be generated using a plurality of computing languages.

[0032] According to a sixth aspect of the present invention there is provided a computing entity comprising:

a data processing means;

a memory means;

an operating system having a set of kernel code, containing a communications protocol stack code;

a plurality of encryption and decryption means;

a first area of said memory means being assigned to said operating system code;

a second area of said memory means being assigned to said plurality of encryption and decryption means, said second memory area being subdivided into a plurality of compartmented memory areas within said second memory area, wherein individual ones of said encryption and decryption means are resident in corresponding respective ones of said plurality of compartmented memory areas.

[0033] Preferably, said computing entity further comprises a director means resident in said first memory area, said director means arranged to input data from and output data to said communications protocol stack.

[0034] Substantially the director means sends a plurality of data streams to said plurality of compartmented memory areas and receives a plurality of data streams from said plurality of compartmented memory areas.

[0035] Said director means may operate to send a corresponding respective data stream to each of said plurality of compartmented memory areas and receive said corresponding respective data stream from said corresponding compartmented memory areas.

[0036] The computing entity may further comprise a plurality of ports resident in said first memory area, said port positioned between a mentioned memory area, said port positioned between a there being at least one said port said compartmented memory area, said port positioned between a said corresponding respective compartmented memory

Brief Description of the Drawings

9 EP 1 079 581 A2

remote computing entity according to a third specific method of the present invention;

[0040] For a better understanding of the invention and to show how the same may be carried into effect, there will now be described by way of example only specific embodiments, methods and processes according to the present invention with reference to the accompanying drawings in which:

Fig. 3 illustrates schematically a plurality of computing entities exchanging data across a communications network, as known in the prior art;

Fig. 2 illustrates schematically protocol stacks for first and second computing entities exchanging encrypted data over a communications network, as is known in the prior art;

Fig. 1 illustrates schematically a fault tolerant entities exchanging data across a first and second layer of said protocol stack in a first memory area;

Fig. 4 illustrates schematically a prior art encryption entities incorporating a prior art encryption system having an Internet protocol security layer and a key exchange protocol for negotiating an exchange of key data;

Fig. 5 illustrates schematically a fault tolerant encryption/decription system for use in transmitting encrypted data between a first and second specific implementation of the present invention;

Fig. 6 illustrates schematically a fault tolerant encryption/decription system for use in transmitting encrypted data between computing entities across a network according to a second specific implementation of the present invention;

Fig. 7 illustrates schematically a plurality of Internet protocol security stacks operating in a plurality of corresponding respective compartment memory areas which are distinct from each other, in an area of user memory according to third specific implementation of the present invention; and

Fig. 14 illustrates schematically a plurality of separate compartment memory areas in an area of user memory distinct from an area of memory reserved for an operating system, the plurality of compartment memory areas each running one or plurality of processes in parallel, and each compartment respective data processing capacity in the form of a corresponding microprocessor, according to fourth specific implementation of the present invention;

Fig. 11 illustrates schematically a flow diagram of steps performed after an Internet protocol security stack on receipt of a data packet from a kernel/operating system according to a fourth specific method of the present invention;

Fig. 12 illustrates schematically a flow diagram of steps performed after an Internet protocol security stack has encrypted or decrypted a data packet according to a fifth specific method of the present invention;

Fig. 13 illustrates schematically a fault tolerant encryption/decription system for use in transmitting encrypted data between computing entities across a network according to a second specific implementation of the present invention;

Fig. 10 illustrates schematically process steps performed by a protocol stack an Internet protocol security stack on receipt of a data packet from a kernel/operating system according to a third specific method of the present invention;

Fig. 9 illustrates schematically a flow diagram summarizing the process steps performed when an application requests to exchange data with a

Fig. 8 illustrates schematically a specific implementation of the present invention showing an Internet protocol security (Internet protocol security) stack located in a user memory space triggering a inter-

net key exchange agent device to negotiate an exchange of keying material with a remote computing entity across a communications network;

Fig. 7 illustrates schematically an overview of process steps implemented by a first communicating entity in response to receiving encrypted data from a second communicating entity across a virtual private network according to a first specific method of the present invention;

Fig. 8 illustrates schematically an overview of process steps implemented by a first communicating entity across a virtual private network according to a second specific method of the present invention;

Fig. 9 illustrates schematically a flow diagram summarizing the process steps performed when an application requests to exchange data with a

Out of the Invention

Detailed Description of the Best Mode for Carrying Out the Invention

[0141] There will now be described by way of example the best mode contemplated by the inventors for carrying out the invention. In the following description numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent however, to one skilled in the art, that the present invention may be practiced without limitation to these specific details. In other instances, well known methods and structures have not been described in detail so as not to unnecessarily obscure the present invention.

[0142] Specific implementation and methods according to the present invention as described herein are aimed at computing entities configured to exchange encrypted data across digital communications networks, including but not limited to local area networks,

wide area networks, the Internet and telecommunications networks.

the Internet protocol security stack due to receipt of malformed data packets. It is likely that there will be pro-

figured to be an interface between the network protocols within the kernel and the Internet protocol suite.

However, it will be understood by those skilled in that the scope of the present invention is as

[0043] Referring to FIG. 5 herein, there is illustrated schematically a method and apparatus according to a best mode implementation of the present invention for robust encryption and decryption of data, for transmission over a communications network. The best mode implementation is described in relation to a virtual pri-

5

[0044] gramming errors within any particular implementation of the internet protocol security protocol.

[0045] According to a preferred embodiment of the present invention, data encryption is effected by employing a standard known encryption protocol — the Internet protocol security protocol, but with the code of

stack 510 running as a process within user memory area 503. The STREAMS standard provides a framework for implementing input/output functions in a modular fashion. This provides an alternative to a traditional UNIX character input/output (IO). Use of the STREAMS framework is described in "STREAMSUXI".

Within the claims herein and not limited by the details of the encryption and decryption algorithms used.

The encrypted or decrypted user data is sent to software port 509 and back to the redirector B. If the user data processed by the internet pro-

[0049] By separating the Internet protocol security stack to run in the second memory area, apart from the memory area used for running the kernel/operating system, additional processing power may be brought bear on the data processing carried out by the Internet protocol security stack 510, as shown schematically in FIG. 5 by the presence of a plurality of additional microprocessors 511, 512.

[0050] Data to be either encrypted or decrypted by the Internet protocol security stack 510 is first transferred from the redirector layer 508 to the Internet protocol security stack 510 in the second memory area, software port 508. The data to be encrypted or decrypted are processed for encryption or decryption.

within user memory. The reduction of data in the operating system/kernel and user memory, and description of data are effected by the processor 500 operating according to a set of instructions stored within the operating system/kernel over a network to a remote computing entity over a network. A encrypted data is passed from reflector layer network interface card 505 and is subsequently reflected over the virtual private network. Encrypted data which has been received by network interface card 510 has been decrypted by network interface card 510 is passed to a protocol stack by reflector layer 508 and is subsequently reflected over a protocol stack by reflector layer 504 which is then ultimately passed to application layer 504 which is within user memory. The reduction of data in the operating system/kernel and user memory, and description of data are effected by the processor 500 operating according to a set of instructions stored within the operating system/kernel

[0049] By separating the Internet protocol security stack to run in the second memory area, apart from first memory area used for running the kernel operating system, additional processing power may be brought bear on the data processing carried out by the Internet protocol security stack 510, as shown schematically in FIG. 5 by the presence of a plurality of additional microprocessors 511, 512.

[0050] Data to be either encrypted or decrypted by the Internet protocol security stack 510 is first transferred from the redirector layer 508 to the Internet protocol security stack 510 in the second memory area, software port 508. The data to be encrypted or decrypted are processed for encryption or decryption within Internet protocol security stack 510. The encrypted or decrypted data are passed back to software port 509 and then back to redirector layer 508.

[0051] Encryption and the complementary decryption of data to be transferred across a virtual private network, according to the best mode presented herein, may be effected using a standard encryption cipher such as the known data encryption standard (DES).

operating system code comprises code for operating network protocols, according to a protocol stack including protocols for the transmission control protocol, including the protocol (TCP/IP), 506, a user datagram protocol (UDP/IP) 507, and a protocol/internet protocol (UDP/IP stack) 507. When data is sent or received through network interface card 505, the data passes up or down the protocol stack comprising the protocol layers 508-507 to be converted and documented between a data form suitable for use by the communication entity and a data format suitable for transmission across a network.

40

Combined within the operating system are the set of protocols configured to process data for transmission across a network. The set of protocols comprises a network protocol stack such as the Transmission Control Protocol/Internet Protocol (TCP/IP) or User Datagram Protocol/Internet Protocol (UDP/IP) stacks. According to a specific implementation of the present invention, the network protocol stack includes the following sub layers:

[1895] (1995) The IBM Reference Manual, Hewlett-Packard

[0049] By separating the Internet protocol security stack to run in the second memory area, apart from first memory area used for running the Kernel/operating system, additional processing power may be brought bear on the data processing carried out by the Internet protocol security stack 510, as shown schematically in FIG. 5 by the presence of a plurality of additional microprocessors 511, 512.

[0050] Data to be either encrypted or decrypted by the Internet protocol security stack 510 is first transferred from the reflector layer 508 to the Internet protocol security stack 510 in the second memory area, software port 508. The data to be encrypted or decrypted are processed for encryption or decryption within Internet protocol security stack 510. The encrypted or decrypted data are passed back to software port 509 and then back to reflector layer 508.

[0051] Encryption and the complementary decryption of data to be transferred across a virtual private network, according to the best mode presented herein, may be effected using a standard encryption cipher such as the known data encryption standard (DES).

The data encryption standard is a block cipher configured to encrypt user data in 64-bit blocks. A 64-bit block of unencrypted data is processed into a 64-bit block of encrypted data. Additionally, DES is a symmetric algorithm requiring a same "key" to both encrypt and decrypt user data. Hence, a same key must be stored securely at both computing entities which are exchange encrypt user data across a virtual private network. The key used for encrypting and decrypting user data with the DES cipher is, in practice, a shared key.

use for transmission over a communications network
accessible via the network interface card 505.
[0044] In prior art systems, the operating system is
partially protected from failure of applications running
user memory. Should an application fail, for any reason,
then whilst the application itself may terminate abruptly
it is unlikely to have any serious consequences for the
operation of the operating system. As described herein
before, a prior art solution to the problem of providing a
software based virtual private network is to insert an
internal protocol security layer within the network proto-
col stack located within the operating system. If, due to
receiving incorrectly formatted data, the internal proto-
col security block crashes because the stack is in a hier-
archical structure of the operating system, then this could crash the
whole operating system. In addition to faults occurring in
the operating system, in addition to faults occurring in

[0049] By separating the Internet protocol security stack to run in the second memory area, apart from first memory area used for running the kernel/operating system, additional processing power may be brought bear on the data processing carried out by the Internet protocol security stack 510, as shown schematically in Fig. 5 by the presence of a plurality of additional mini-processors 511, 512.

[0050] Data to be either encrypted or decrypted by the Internet protocol security stack 510 is first transferred from the reflector layer 508 to the Internet protocol security stack 510 in the second memory area, software port 508. The data to be encrypted or decrypted are processed for encryption or decryption within Internet protocol security stack 510. The encrypted or decrypted data are passed back to software port 509 and then back to reflector layer 508.

[0051] Encryption and the complementary decryption of data to be transferred across a virtual private network, according to the best mode presented herein, may be effected using a standard encryption cipher such as the known data encryption standard (DES). The data encryption standard is a block cipher used to encrypt user data in 64-bit blocks. A 64-bit block of unencrypted data is processed into a 64-bit block of encrypted data. Additionally, DES is a symmetric algorithm requiring a same "key" to both encrypt and decrypt user data. Hence, a same key must be stored securely at both computing entities which are exchange encrypted data across a virtual private network. The key used for encrypting and decrypting data with the DES cipher is typically 8 characters long. A number of private keys stored in a key database are associated with the Internet protocol security stack 510 in a range one to several thousand, the number of private keys stored being related to a number of connections required to be made between a first computing entity and a plurality of remote computing entities. In a way of example, a digital computer configured as a server may typically require twice as many private keys associated with a single Internet protocol security stack as the number of connections to client machines.

[0052] In an alternative embodiment of the present invention, user data to be transmitted across a virtual private network may be encrypted and decrypted using a known cipher such as Triple-DES. This cipher provides enhanced resistance to cryptanalysis at the cost of requiring a private key, typically 24 characters

to a variety of operating systems such as Win-NT4.0. It is preferable that such methods as a remote protocol security stack 510 be passed up the protocol stack by reflector layer 505 and is subsequently routed to an application 504 which is a set of data within user memory. The redirection of data within the operating system/kernel and user memory, and which has been subsequently decrypted by the network interface card 505 and is passed up the protocol stack by reflector layer 505 and is subsequently routed to a remote computing entity which has been encrypted for transmission over a network to a remote computing entity in a remote computing entity within a remote digital computer or other computing system such as HPUX10.24.

Referring to Fig. 8 herein there is illustrated invention. As described herein before, a radio-frequency transceiver 508 within the protocol stack logically within first memory area 502 reserved for the application 500 and 501 may be considered completely as being located within a single memory compartment 503 which is a further logical sub division of memory area 503. User applications and processing applications 600 and 601 may be considered completely as being located within a single memory compartment 602 and 603 respectively as indicated by their corresponding respective security labels as described hereinbefore.

If application 600 requires to transfer data to a network protocol stack located within the first memory compartment 503 then the data packetized and reflected by reflector layer 505 and reflected and reflected by the network protocol stack to the software port 506 data packet to be transmitted is sent to the network protocol stack located within the first memory compartment 503. Every data packet transmitted must, in order to conform with the protocol security protocol, be checked by the

Internet protocol security stack 510 against a security policy database associated with key database 602.

Data describing a security policy, comprising a predetermined set of rules for dealing with data packets, is stored in a security policy database 605 and is accessible by the internet protocol security stack 510. If internet protocol security stack 510 detects that it has not received a security association for transforming a particular type of data, for example, to a remote host then implemented as a process running outside of the operating system and is resident in user memory, the term "security association" as used herein refers to a plurality of security services provided for and attached to individual IP packets securely transmitted across a virtual private network. Each security association inherits a same security label as a corresponding respective data packet with which it is associated. A security association is uniquely identified by:

- A security parameter index (SPI);
- An IP destination address; and
- A security protocol identifier.

[0056] Referring to Fig. 7 herein, there are illustrated schematically in overview, a set of process steps implemented by a first computing entity upon receipt of encrypted data from a second computing entity across a virtual private network according to the best mode described herein. In step 710, a memory means within the first computing entity is split into at least a first area and a second area. The first memory area is configurable to store the operating system and the second memory area is configurable to act as user memory to store user applications and the like. In step 720, encrypted IP data packets, encrypted according to the internet security protocol, are received from a second computing entity across a virtual private network via network interface card 505 and the encrypted data is temporarily stored within the network protocol stack located within the first memory area. In step 730, the encrypted data is transferred up through subsequent layers of the network protocol stack until they enter redirector layer 508. When the encrypted data enters the redirector layer 508 they are redirected to internet protocol security stack 510 which is located within the second memory area 503 via a port 509. In step 740, the encrypted data stored within the internet protocol security stack 510 decrypts the encrypted data. In step 750, decrypted data is passed from internet protocol security stack 510 via software port 509 back to the redirector layer 508 within the network protocol stack located within the first memory area.

[0057] Referring to Fig. 8 herein there are illustrated

schematically an overview of process steps implemented during the transfer of user data from the first computing entity to a second computing entity across a virtual private network according to the best mode presented herein.

5 The first computing entity is split into at least a first area 502 and a second area 503. The first memory area stores the operating system and the second memory area is configured to act as user memory to store user applications and the like. In step 810, a user application 10 wishing to send data to a second computing entity across a virtual private network writes unencrypted data from the second memory area (which is user memory) into the network protocol stack located within the first memory area. In step 820, the unencrypted data passes down the network protocol stack until it enters redirector layer 508. Redirector layer 508 writes the data from the network protocol stack to the internet protocol security stack located in the second memory area 503 via port 509. In step 830, the internet protocol security stack 510 encrypts the data. In step 840, the encrypted data is written from the internet protocol security stack within the second memory area back into redirector layer 508 of the network protocol stack located within the first memory area via port 509.

- A security parameter index (SPI);
- An IP destination address; and
- A security protocol identifier.

[0058] Referring to Fig. 9 herein, there are illustrated schematically process steps implemented by applications within user memory and the operating system according to the best mode described herein. In step 900 an application 801 requires to transmit data to a remote host. The application, which may comprise a plurality of processes operating in user memory, creates a socket buffer within the operating system. A user application creates a socket by issuing a system call from the operating system. Having received a system call from an application 801, the operating system creates a kernel via port 509. Alternatively, if the received data packet does not require decrypting then, in step 1020, the unprocessed data packet is forwarded to the kernel via port 505. If, for example, the data packet received from the operating system is inspected by the operating system to act as an intermediary in a transfer of data between the protocol stack and the application, in step 1010, the application writes data to the socket buffer (also known herein as a network connection). Each data packet written to the socket buffer inherits a security label having a same level of security as the application "banned" remote hosts are dropped. The process of dropping a data packet as described herein comprises overwriting the location within the memory containing the dropped data packet with another data packet.

[0059] Fig. 10 illustrates schematically process steps applied to outgoing data packets generated, for example, a process operating within the user memory of the local host. If a user process 800 wishes to transmit a data packet or plurality of data packets to a remote host across a network then, in step 1000, a data packet to be transmitted enters the redirector layer 508 within the protocol stack and is redirected via port 509 to the internet protocol security stack residing within the compartment logically located within the user memory area 503. In step 1010, the security database associated with key database 602 is consulted to determine whether the data packet received from user process 600 is to be encrypted prior to transmission across the network. If it is determined that the received data packet is to be encrypted then in step 1030, the data packet is encrypted and, in step 1040, returned to the protocol stack via port 509. If the security policy database specifies that the data packet to be transmitted is not to be encrypted then in step 1020 the unprocessed packet is returned via port 509 to the protocol stack. Any packets banned from transmission by the local computer

schematically process steps implemented by internet protocol security stack 510 in response to receipt of an IP packet from the operating system. In step 1050, the internet protocol security stack receives a data packet from a port logically located within the operating system. In step 1010, a security policy database (SPD) associated with key database 602 is consulted by the internet protocol security stack to determine how the received data packet should be treated. For each data packet received by the internet protocol security stack there are three possible outcomes of this consultation process as follows:

- "Apply internet protocol security";
- "Bypass internet protocol security";
- "Drop".

[0060] For incoming data packets received from the remote host across a LANMAN each packet received from the operating system is inspected by the protocol security descriptor. If necessary by examining a security descriptor data comprising a part of a security association, data logically associated with the data packet, in step 1030, decryption of the encrypted data packet is performed and, if successful, the decrypted data packet is sent back to the TCP/IP stack within the kernel via port 509. Alternatively, if the received data packet does not require decrypting then, in step 1020, the unprocessed data packet is forwarded to the kernel via port 505. If, for example, the data packet received from the remote host from which recipient of data is forbidden, then any data packets originating from the "banned" remote hosts are dropped. The process of dropping a data packet as described herein comprises overwriting the location within the memory containing the dropped data packet with another data packet.

[0061] Fig. 10 illustrates schematically process steps applied to outgoing data packets generated, for example, a process operating within the user memory of the local host. If a user process 800 wishes to transmit a data packet or plurality of data packets to a remote host across a network then, in step 1000, a data packet to be transmitted enters the redirector layer 508 within the protocol stack and is redirected via port 509 to the internet protocol security stack residing within the compartment logically located within the user memory area 503. In step 1010, the security database associated with key database 602 is consulted to determine whether the data packet received from user process 600 is to be encrypted prior to transmission across the network. If it is determined that the received data packet is to be encrypted then in step 1030, the data packet is encrypted and, in step 1040, returned to the protocol stack via port 509. If the security policy database specifies that the data packet to be transmitted is not to be encrypted then in step 1020 the unprocessed packet is returned via port 509 to the protocol stack. Any packets banned from transmission by the local computer

ising entity, after consultation of the security policy database, are dropped.

[0062] Referring to Fig. 11 herein, there are illustrated schematically process steps performed by the internet protocol security stack receives a data packet from a port logically located within the network protocol stack according to the best mode presented herein. In step 1100, a security labeled data packet is received from an internet protocol security stack via a port. The received data packet may be encrypted if, for example, the data packet has been originated by a local application located in a same compartment in user space as the internet protocol security stack performing the encryption. Alternatively, the data packet may be decrypted in the case of a incoming data packet received from a remote computing entity across a communications network. In step 1110, having received a data packet from an internet protocol security stack the redirector layer 508 reflects the data packet from the port into the protocol stack within the kernel/operating system. In step 1140, the encrypted data packet is sent down the protocol stack to a network adapter card prior to transmission across the LANMAN to a remote computing entity. The remote computing entity across the LANMAN may comprise a physical network such as an Ethernet or Asynchronous transfer mode connection or the like.

[0063] In summary, Figs. 7, 8, 9, 10 and 11 represent the flow of data packets from a user application in user memory down through the network protocol stack until being redirected to a corresponding internet protocol security stack in a same compartment of user memory as the user application originating the data. Additionally, these figures illustrate the flow of data packets from an internet protocol security stack into the redirector layer prior to transmission across a LANMAN. Figs. 7, 8, 9 and 11 also represent the flows of data packets from a computer across a communications network and entering the network protocol stack of the local computing entity via a network interface card before being redirected into an internet protocol security stack located within user memory. After decryption, the data packets are returned to the network protocol stack prior to the data packets proceeding up the network protocol stack for further processing by the kernel and/or user applications within user memory.

[0064] Locating an internet protocol security stack outside of an area of memory containing the operating system offers significant advantages over prior art

encryption systems used in virtual private networks. In particular, location of an internet protocol security stack in user memory rather than within the operating system may render a computer using such an implementation

more tolerant to programming errors arising from the internet protocol security stack and more able to deal with malformed data packets. Prior art secure networks implementing the internet protocol security must necessarily be using comparatively recently written computer programs and the internet protocol security itself has been introduced only comparatively recently. Hence, any computer code implementing the internet protocol security will, inevitably, have incorporated a larger number of errors than the more well established computer code such as operating within the operating system incorporation of the code associated with an internet protocol security stack into the network protocol stack as implemented in prior art virtual private networks renders the entire operating system open to failure should the comparatively untested internet protocol security stack fail.

[0085] In the present best mode, incorporation of an internet protocol security stack into user memory effectively isolates the operating system and in particular the kernel from internet protocol security stack failure thereby rendering the system more tolerant to faults.

Additionally, implementation of an internet protocol security stack within user space as described herein allows for the possibility of parallel processing of data packets by a plurality of processors enabling a potential increase in the throughput of data transmitted or received across virtual private network. It has been found experimentally by the inventor that the steps of redirecting data packets from a network protocol stack into user memory and then returning the data packets to the redirector layer does add a small time penalty compared with data packets passing straight through the protocol stack without redirection. Typically, it is found that redirection of data packets into user space slows the transfer of data packets by approximately 4% compared with un-redirected packets. However, the potential substantial gains in processing speed resulting from the use of multiple processors to perform the internet protocol security processing within the internet protocol security stack is located in user memory significantly outweighs the time penalties involved in data packet redirection.

[0086] Referring to Fig. 12 herein there is illustrated schematically a second specific implementation of the present invention, located within an area of user memory a single memory area component 1250 is assigned to a plurality of internet protocol security stacks illustrated herein by way of example by first and second internal protocol security stacks 1210 and 1220.

50 In order to minimize the possibility of experiencing internet protocol security stack failure which can result in sending corrupted data packets back to the operating system it is possible to implement several versions of the internet protocol security stack within a single memory area 1250 which is assigned to use by the internet protocol security stacks and is separate from a memory area used by the kernel of the operating system. Each

implementation of the internet protocol security stack may be generated using a different compiler or, for example, a different programming language. Whilst it is possible that each implementation of the internet protocol security stack will contain programming errors there is a very low probability that different implementations of the internet protocol security stack will manifest the same errors in response to an identical input. According to a second embodiment of the present invention, it is arranged that a plurality of internet protocol security stacks 1210, 1220 are provided within a single memory area each internet protocol security stack accessing a corresponding respective key database 1260, 1270. The data stored within the plurality of key databases 1260 and 1270 are identical or overlapping. Comparison of the outputs of the plurality of redundant internet protocol security stacks can be used to identify a correct output. According to the second embodiment of the present invention the plurality of redundant internet protocol security stacks 1210, 1220 within a single compartment all exchange data with a redirector layer 508 within a network protocol stack via a single port 220.

[0087] Locating a plurality of different implementations of an internet protocol security stack within a single compartment of memory area, each internet protocol security stack accessing a corresponding respective key database each key database containing identical or substantially similar security data yields the following advantages:

30 • Increased fault tolerance by isolating an internet protocol security stack in user memory logically distinct from an area of memory containing an operating system;

35 • locating multiple redundant internet protocol security stacks accessing the same security database information reduces the probability of a bug within a single internet protocol security stack sending corrupted data back to the network protocol stack again yielding improved fault tolerance;

40 • locating a plurality of internet protocol security stacks in user memory enables the processing power of a plurality of processors to be used in encryption and decryption thereby yielding improved data throughput rates.

[0088] Experiments have been carried out to test the throughput rates for implementations of the internet protocol security stacks described herein above comprising one internet protocol security stack per compartment in user space 45 against two internet protocol security stacks per compartment, the throughput rates for implementations of the internet protocol security stacks, from faults and crashes caused by other ones of the plurality of internet protocol security stacks. Each internet protocol security stack operates as a separate processor running in its corresponding respective compartment memory area. If an internet protocol security processor crashes, basic operation of the kernel of the operating system remains unaffected, and the effect on the kernel is limited to the port 50 directed by redirector layer 1403 to the corresponding respective compartment memory area 1408-1410. A first stream of

the results were that the second implementation of the present invention yields higher data throughput rates by distributing the processing of encrypted data packets between multiple internet protocol security stacks and a plurality of processors. For the case of one internet protocol security stack per compartment, an average throughput of 321 kbytes was achieved. For a corresponding second implementation having two internet protocol security stacks per compartment, an average throughput of 483 kbytes was achieved.

[0070] Referring to Fig. 13 herein, there is illustrated schematically a third specific implementation according to the present invention, in which is provided a second area of user memory 1300, which is distinct from a first memory area 1301 containing a kernel code 1302 of an operating system within which is a protocol stack 1303 containing a redirector layer 1304 and a plurality of ports 1305, 1306, for example STREAMS ports, and a plurality of further memory areas 1307, 1308 resident within the area of user memory 1300, the user memory area being an area reserved for applications and processes running under control of the operating system of the computer entity, in which the plurality of compartments of memory areas within the user memory each contain at least one corresponding respective internet protocol security stack 1309, 1310, having a respective key database 1311, 1312, corresponding respectively to the first and second areas, such that a plurality of internet protocol security processes run in parallel and simultaneously within the user memory, each internet protocol security stack running in its own separately assigned compartment memory area, which is logically isolated from and logically distinct from other compartment memory areas.

[0071] The compartment memory areas preferably are each assigned specifically to one or more internet protocol security stacks and their associated key databases, but optionally may also run other processes, such as a plurality of user applications 1313, 1314 which interface with the operating system in the first memory area 1301, without passing through the STREAMS ports.

[0072] Isolation of separate internet protocol security stacks in separate compartment memory areas may achieve increased robustness and reliability of the operating system and its kernel through isolation of the internet protocol security stack from the first memory area containing the code of the kernel, and additionally may provide increased robustness for each of the plurality of internet protocol security stacks, from faults and

50 crashes caused by other ones of the plurality of internet protocol security stacks. Each internet protocol security stack operates as a separate processor running in its corresponding respective compartment memory area. If an internet protocol security processor crashes, basic operation of the kernel of the operating system remains unaffected, and the effect on the kernel is limited to the port directed by redirector layer 1403 to the plurality of ports 1404-1406 and to the corresponding respective compartment memory areas 1408-1410. A first stream of

55 data running in a single compartment memory area is running in a same compartment memory area as an internet protocol security stack which has crashed, then those processes are also vulnerable to fault caused by the crash of the internet protocol security stack in that particular compartment memory area. Whilst running of separate processes in a same compartment memory area assigned to an internet protocol security stack is possible in implementations of the invention, in the best mode, a compartment memory area is assigned specifically and exclusively to an internet protocol security stack, and its associated key database, so that a fault in an internet protocol security stack in a given compartment does not automatically cause the processing of other internet protocol security stacks in other compartments, or the kernel, or other processes such as applications being run to fail. Providing different internet protocol security stacks in different compartments of memory areas also has an advantage that different streams of data can be processed independently by their own dedicated internet protocol security stacks, independently of other data streams.

[0073] Referring to Fig. 14 herein, there is illustrated schematically a fourth specific implementation according to the present invention, in which a first memory area 1400 of a first computing entity contains an operating system kernel 1401 having a transport protocol stack 1402 which includes a redirector layer 1403, there being provided a plurality of STREAMS ports 1404-1408 within first memory area 1400 for communicating one or more packet data streams to a plurality of internet protocol security stacks operating in an area of user memory 1407 which is assigned as a plurality of compartment memory areas 1408-1410, each compartment memory area having the benefit of processing capabilities provided by a corresponding respective micro-processor 1411-1413. In Fig. 14, processes carried out by the internet protocol security stacks are illustrated as circles, one circle per separate process running. A single internet protocol security stack may process a set of separate individual packet data streams on a multiplex basis, or alternatively a plurality of different internet protocol security stacks in different compartment memory areas may be applied, each having its own processing power supplied by a corresponding micro-processor to process a single stream of data packets in parallel.

[0075] A plurality of packet data streams passing down protocol stack 1402 in the kernel 1401 are directed by redirector layer 1403 to the plurality of ports 1404-1406 and to the corresponding respective compartment memory areas 1408-1410. A first stream of

packet data may be routed from the redirector layer 1403 through a STREAMS port 1404, into a first compartment 1403, and the packet data stream may be encrypted processed, either to encrypt the stream or decrypt the stream in a plurality or processes shown schematically as circles in FIG. 14, with the processing power being supplied by a first microprocessor 1411. The one/plus one processed line data stream is then burned through the first port 1404 back to the redirector layer 1403 and on to the next layer in the communications protocol stack 1402 in encryption processed format.

[0076] Similarly, a second data stream passing up or down the communications protocol stack 1402 is passed by redirector layer 1403 to a second port 1405, which passes the data stream out of the first memory area 1400 into the user memory area 1407, and in particular into a second compartment memory area 1409 where the second packet data stream is encryption processed, represented as a circle in FIG. 14, with processing power for the encryption processing being provided by a second microprocessor 1412. The encryption processed second data stream is returned via the second STREAMS port 1405 back to the redirector layer 1403 and from the redirector layer onto the next layer either above or below, depending on which way the data packet stream is passing through the stack.

[0077] Providing a plurality of different processes running in a plurality of compartment memory areas which are distinct from each other may achieve the advantages mentioned herein before of robustness in that if one process fails in a first compartment memory area other processes in other compartment memory areas may continue uninterrupted. By applying separate processing capability in the form of a corresponding respective microprocessor to each separate compartment memory area, as well as enhancing reliability and robustness, the amount of throughput of data processing in terms of Mbytes/s can be enhanced using the architecture and method of operation illustrated schematically in FIG. 14 herein.

Claims

1. A method of processing encryption data in a communications protocol stack 1402, comprising the steps of: assigning a memory means of said computing entity into a plurality of memory areas (740); receiving encrypted data (720); storing said encrypted data in a first memory area (740); and outputting said data to said first memory area (740);
2. A method as claimed in claim 1, wherein said first memory area is logically distinct from said second memory area.
3. A method as claimed in claim 1 or 2, wherein said first memory area is configured to contain code of a kernel of said operating system.
4. The method as claimed in claim 1 or 2, wherein said second memory area is not used for storage of code of a kernel of said operating system.
5. The method as claimed in claim 1, wherein said step of decrypting said encrypted data stored in said second memory area is logically distinct from said first memory area.
6. A method of processing encryption data in a communications protocol stack 1402, comprising the steps of: assigning a memory means of said computing entity into a plurality of memory areas (810); storing decrypted data in a first said memory area of said plurality of memory areas, said first memory area being assigned for use by a kernel code of an operating system of said first computing entity (820); receiving an encrypted data packet; processing said data packet according to at least one packetization protocol of said operating system in said first memory area; outputting said data packet to said second memory area;
7. A method as claimed in claim 6, wherein said second memory area is logically distinct from said first memory area.
8. A method as claimed in claim 6 or 7 wherein said

putting entity:

writing said encrypted data stored in said first memory area into a second memory area associated with said computing entity (730);

decrypting said encrypted data stored in said second memory area (740); and

writing said decrypted data from said second memory area to said first memory area (750).

first memory area is configured to contain code of said operating system.

first memory area (500) and a second memory area (520) and characterized by further comprising:

9. A method as claimed in claim 8 or 7, wherein said second memory area is not used for storage of code of a kernel of said operating system.

10. A method as claimed in claim 8, further comprising the step of redirecting said encrypted data from said operating system in said first memory area to an encryption/decryption stack resident in said second memory area.

15. A digital computer configurable for receiving digital data transmitted across a communications network said digital computer comprising:

11. A method as claimed in claim 9, comprising the step of directing said data to said second memory area from said first memory area.

16. A digital computer configurable for writing data from said first memory area (500) to said second memory area to said first memory area, and encryption means (509) logically located within said second memory area, said encryption means configurable for encrypting said digital data.

12. A method as claimed in claim 6, wherein said step of encrypting said data stored in said second memory area is carried out by an internet protocol security program stored in said second memory area.

17. A digital computer as claimed in claim 15, wherein said second memory area (520) for decrypting data stored in said second memory area (520), said decryption means (509) being configurable for decrypting said data.

13. A method of processing encrypted data in a communications protocol stack 1402, comprising the steps of: assigning a memory means of said computing entity into a plurality of memory areas (810);

a processor; and

a memory means,

wherein said memory means is divided into first and second memory areas, wherein said first memory area contains code of an operating system of said computer entity, said

method comprising the steps of:

receiving an encrypted data packet;

processing said data packet according to at least one packetization protocol of said operating system in said first memory area;

outputting said data packet to said second memory area;

outputting said data packet to said second memory area;

processing said data packet according to a processing algorithm in said second memory area, and

returning said processed data packet to said operating system in said first memory area.

14. A digital computer configurable for transmitting digital data across a communications network, said digital computer comprising:

15. A digital computer as claimed in claim 15 or 16, wherein said encryption means logically located within said second memory area comprises:

16. A digital computer as claimed in claim 15 or 16, wherein said encryption means logically located within said second memory area comprises:

a plurality of internet protocol security stacks (1210, 1220); and
 a plurality of data bases (1280, 1270) configurable to contain key data for encrypting said data, wherein each data base of said plurality of databases contains a same key data.

19. A digital computer as claimed in claim 18, wherein said plurality of internet protocol security stacks are generated using a plurality of computing languages.

20. A computing entity comprising:
 a data processing means;
 a memory means;

an operating system having a set of kernel code, containing a communications protocol stack code;

a plurality of encryption and decryption means assigned to said operating system code;

a second area of said memory means being assigned to said plurality of encryption and decryption means, said second memory area being sub-assigned into a plurality of compartmented memory areas, wherein individual ones of said encryption and decryption means are resident in corresponding respective ones of said plurality of compartmented memory areas.

21. The computing entity as claimed in claim 20, further comprising a director means resident in said first memory area, said director means arranged to input data from and output data to said communications protocol stack.

22. The computing entity as claimed in claim 21, wherein said director means sends a plurality of data streams to said plurality of compartmented memory areas and receives a plurality of data streams from said plurality of compartmented memory areas.

23. The computing entity as claimed in claim 21, wherein said director means operates to send a corresponding respective data stream to each of said plurality of compartmented memory areas, and receive said corresponding respective data stream from said corresponding compartmented memory areas.

ond compartmented memory area is assigned to said second process, and said first memory area is assigned to an operating system of said computing entity.

27. The method as claimed in claim 26, comprising the step of running a plurality of processes in a said compartmented memory area, for processing a single said packet data stream.

28. The computing entity as claimed in claim 20, comprising a plurality of ports resident in said first memory area, there being at least one said port positioned between a said corresponding respective compartmented memory area, and a director means for directing data streams to and from said plurality of compartmented memory areas, said director means being resident in said first memory area.

25. The computing entity as claimed in claim 20, comprising a plurality of data processor means, said plurality of data processors providing processing capability for carrying out data processing operations within said plurality of compartmented memory areas.

26. A method of encryption processing a plurality of packet data streams between first and second layers of a communications protocol stack, said method comprising the steps of:

receiving a first said data packet stream from a first layer of said protocol stack in a first memory area;

sending said data packet stream to a first compartmented memory area;

running an encryption process on said first data packet stream in said first compartmented memory area for encryption or decryption of said data packet stream;

returning said processed data packet stream from said first compartmented memory area to a second layer of said communications protocol stack in said first memory area;

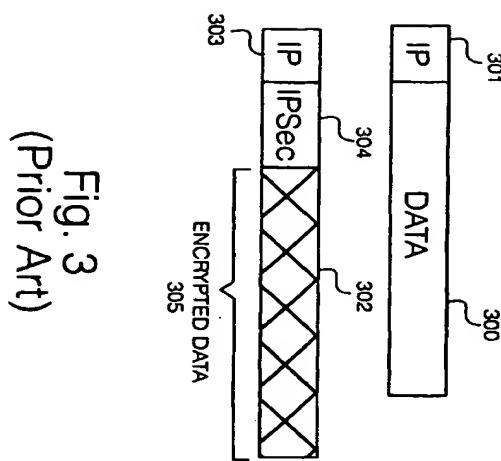
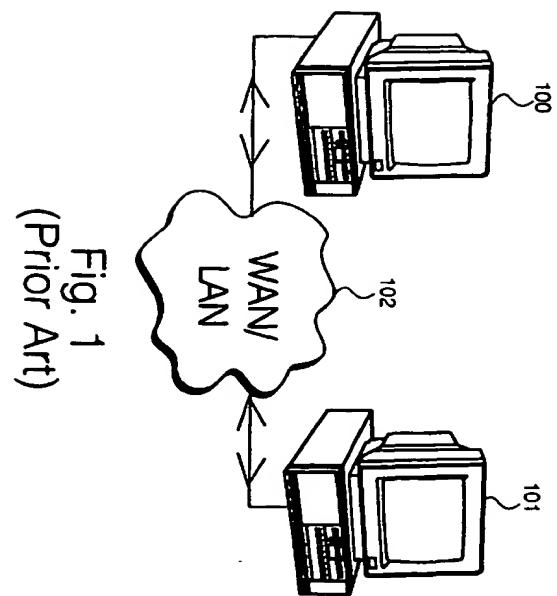
receiving a second packet data stream from a said first or second layer of said communications protocol stack in said first memory area;

sending said second data packet stream to a second compartmented memory area;

encryption processing said second data packet stream in said second compartmented memory area for encryption or decryption of said data packet stream; and

returning said processed second packet data stream to the other one of said first or second layers of said communications protocol stack in said first memory area,

wherein said first compartmented memory area is assigned to said first process, said sec-



17

18

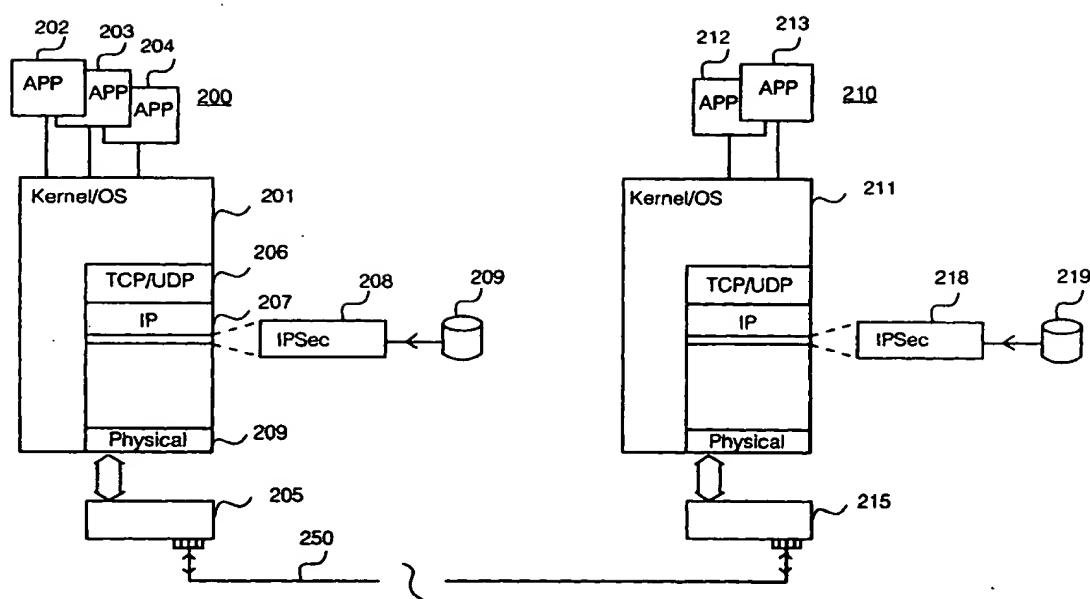


Fig. 4
(Prior Art)

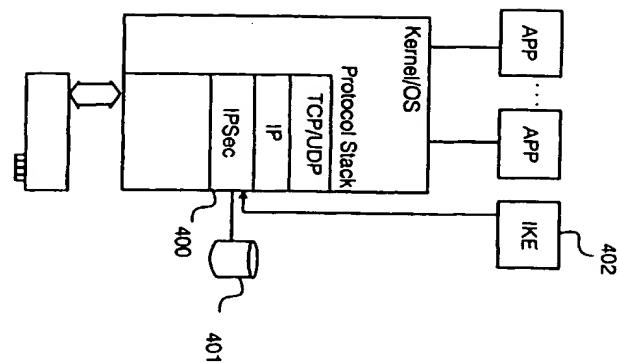
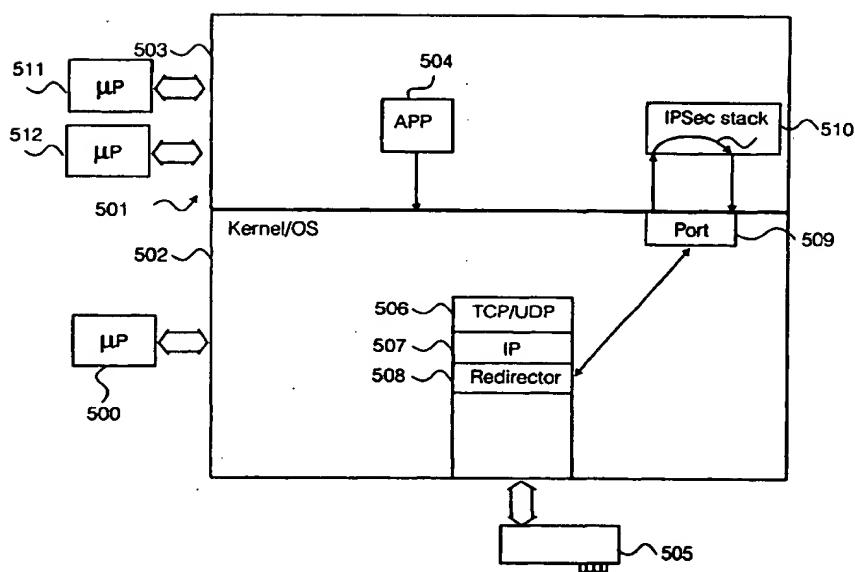


Fig. 5



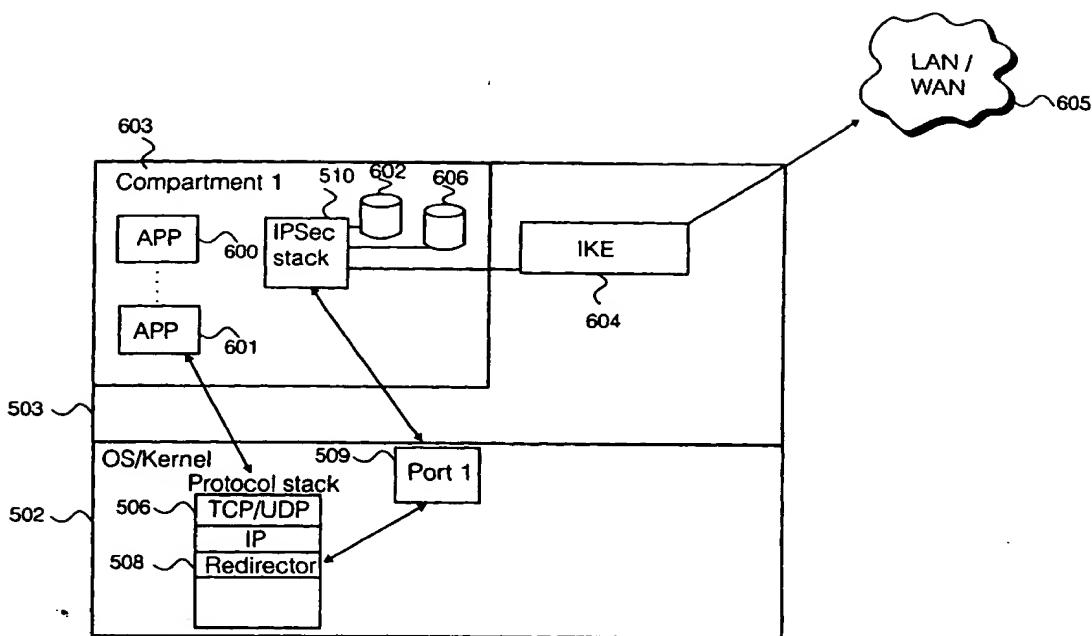
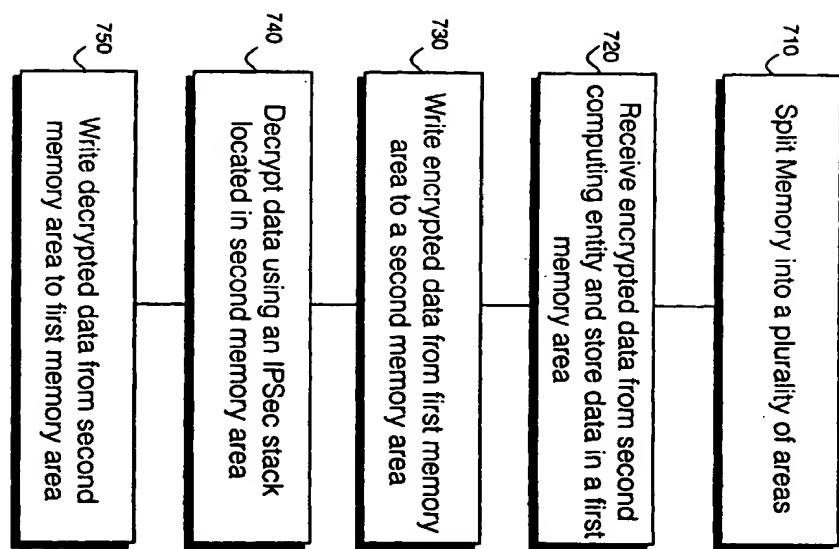


Fig. 6

Fig. 7



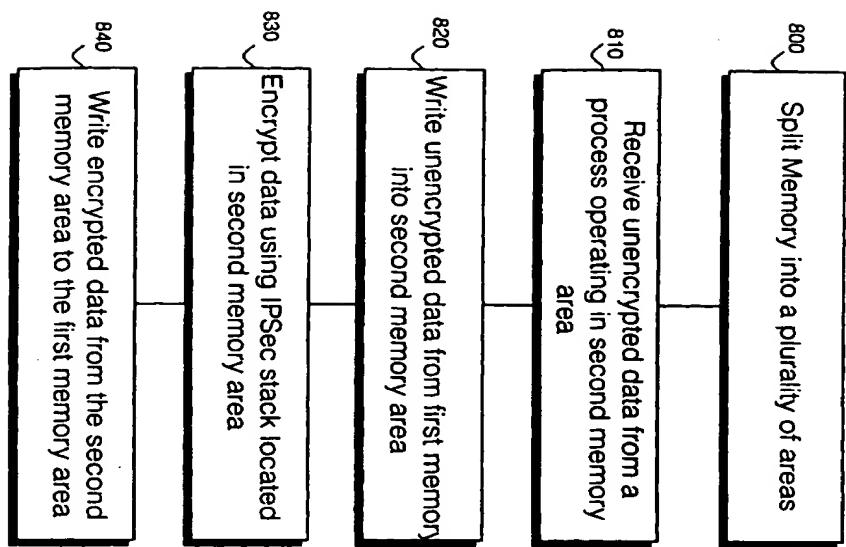


Fig. 8

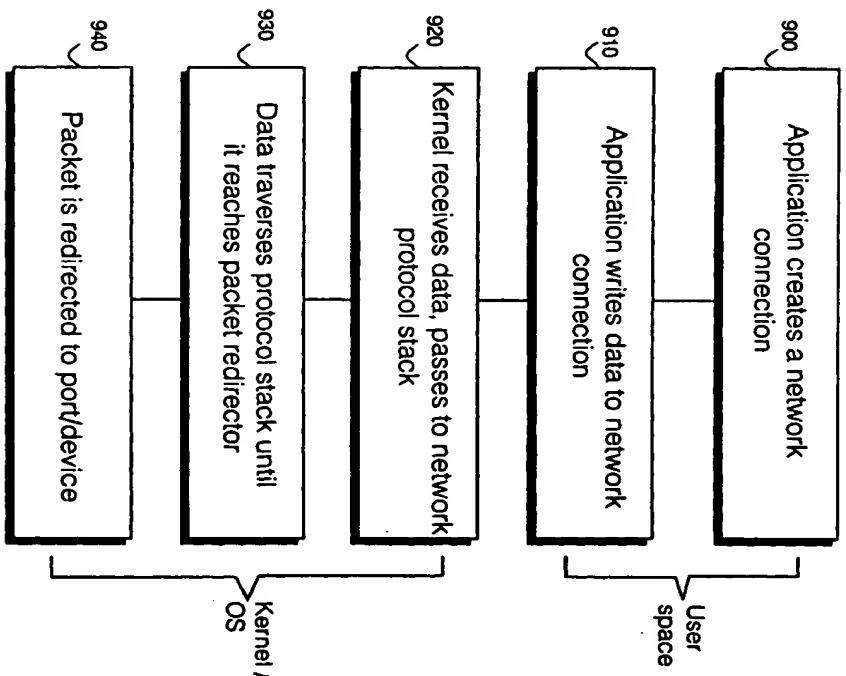


Fig. 9

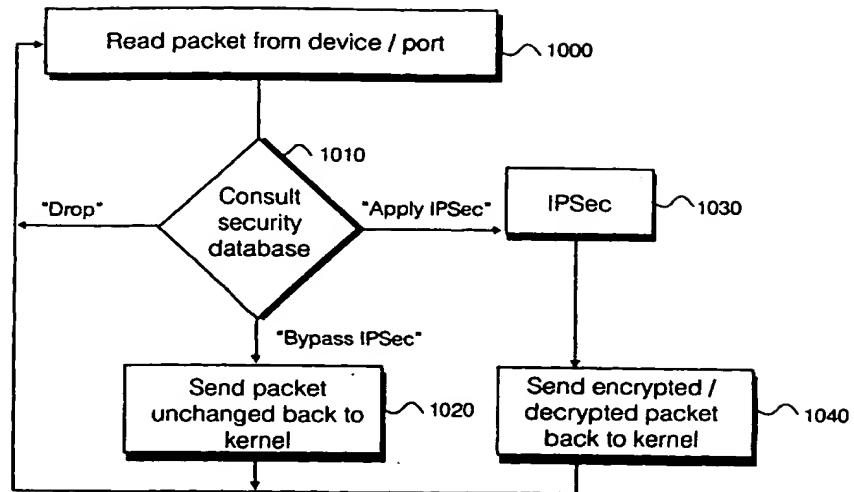


Fig.10

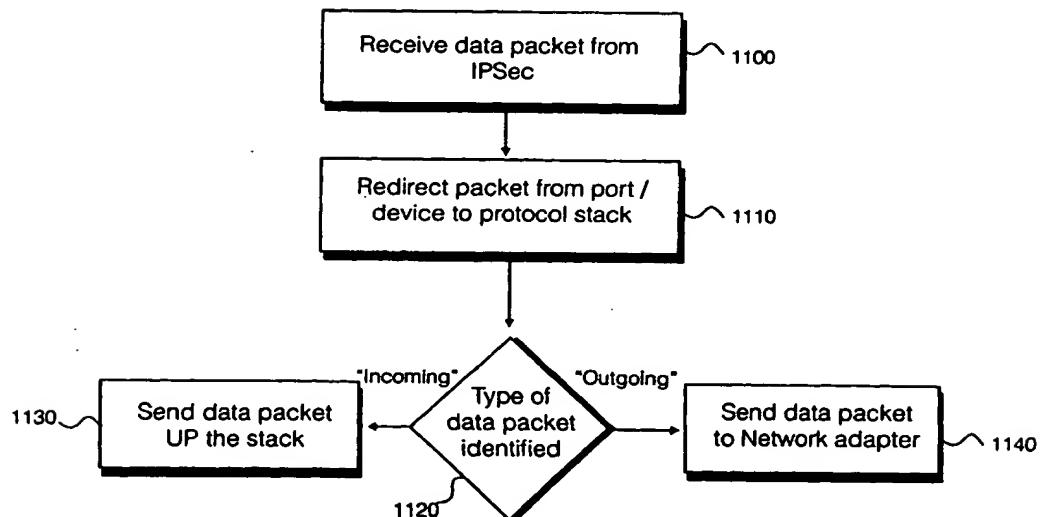


Fig. 11

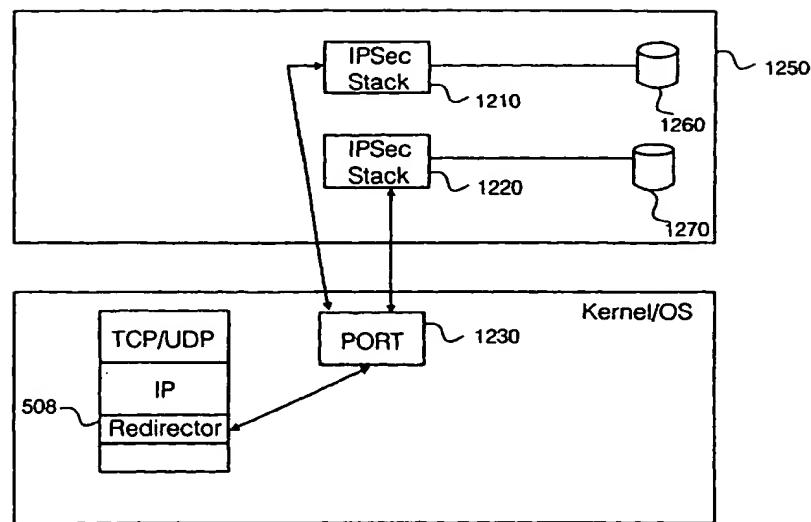


Fig. 12

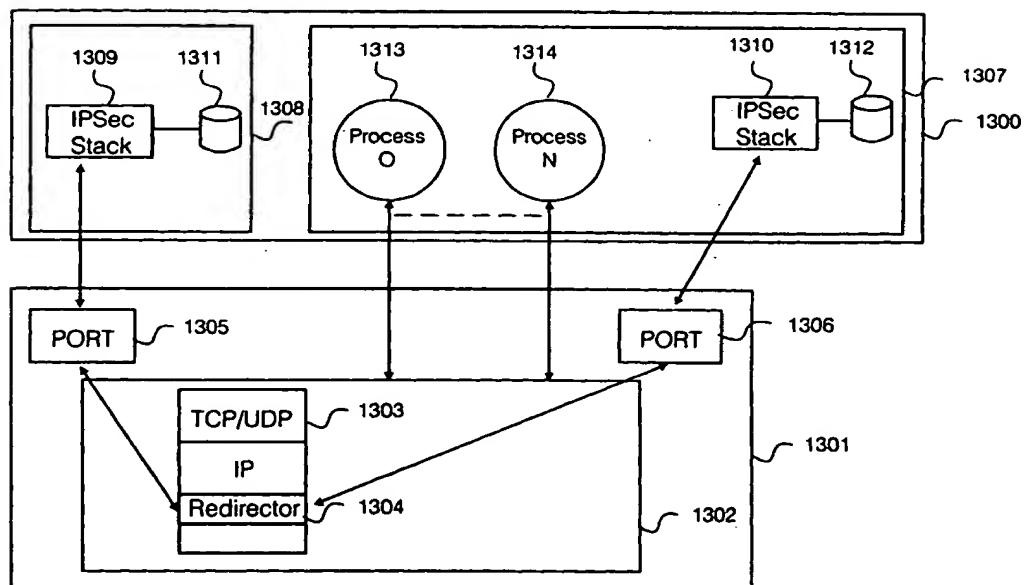


Fig. 13

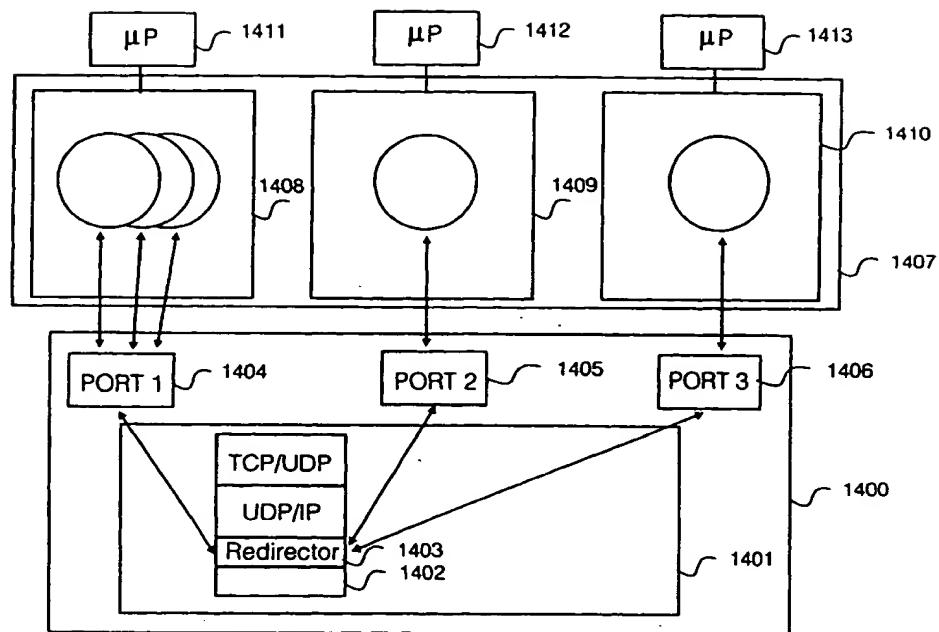


Fig. 14